

# **Ensuring Security and Compliance of Your EMC Documentum Enterprise Content Management System: A Collaborative Effort of EMC Documentum and RSA**

*Applied Technology*

---

***Abstract***

This white paper discusses the challenges with securing content at rest and in motion and presents considerations to assist customers in developing a comprehensive Secure Content Management strategy that will satisfy corporate information security goals and requirements. The core security capabilities of the EMC® Documentum® Content Management System augmented with RSA's security technology provide the technical foundation of this strategy.

November 2009

---

---

Copyright © 2009 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

All other trademarks used herein are the property of their respective owners.

Part Number h4753

---

## Table of Contents

<b>Executive summary .....</b>	<b>4</b>
<b>Introduction .....</b>	<b>4</b>
Audience .....	4
<b>Defining Secure Content Management .....</b>	<b>4</b>
Securing access to the content management system .....	4
Securing content at rest .....	5
Securing content in motion and in use.....	5
Monitoring and reporting on compliance of the environment.....	5
<b>Securing CMS access.....</b>	<b>6</b>
Integrating with the enterprise identity management infrastructure .....	6
Enforcing strong authentication and enabling single sign-on .....	6
<b>Securing content at rest.....</b>	<b>6</b>
Discovering sensitive content .....	6
Controlling access to sensitive content.....	6
Ensuring confidentiality of sensitive data .....	7
Repository encryption.....	7
Digital shredding.....	8
Ensuring authenticity and integrity of sensitive data.....	8
<b>Security content in motion and in use .....</b>	<b>8</b>
Enabling secure communication .....	8
Enforcing sensitive content policies .....	8
Information Rights Management.....	9
<b>Compliance .....</b>	<b>10</b>
Logging in the content management system .....	10
Monitoring and reporting security events .....	11
<b>Conclusion .....</b>	<b>11</b>
<b>References .....</b>	<b>11</b>

---

## Executive summary

During the past several decades, the volume of digital content created and managed by companies has grown, and continues to grow, explosively. This content often contains proprietary information and trade secrets; sensitive employee, customer and partner data; confidential executive communications; corporate financial information; and personal health information. Unlike structured data stored in secured databases, documents - and especially e-mails - are much more mobile. Documents are e-mailed to co-workers, partners, and customers; shared on file systems and Web servers; and, frequently, printed and tossed in the recycling bin.

To deal with the ever-increasing risks of this information being used inappropriately, stolen, or inadvertently leaked, businesses have taken content security very seriously. Securing documents has become a major issue and challenge for companies and organizations, especially as government regulations impose additional burdens for managing, sharing, archiving, and discovering this information. A single lapse in content protection can result in very large fines, lost revenue and customers, legal expenses, embarrassment, and even jail time for the involved executives. These trends toward increasing digital content and increasing regulations show no sign of slowing.

Enterprise content management (ECM) offered a good first step in managing "content at rest" information security. While in the repository, sensitive documents are access controlled and protected. Audit logs capture who accessed the information and when. Unfortunately, content locked away in a repository is not very useful. "Content in motion" is at significantly greater risk as it is moved within an organization and shared with customers and business partners. Technologies like document encryption, Information Rights Management (IRM), and Data Loss Protection (DLP) address different use cases and security requirements.

A comprehensive Secure Content Management solution must provide the appropriate level of protection, at the right time, while simultaneously being practical in today's fast-paced business environment. Protecting all types of content is not the same as protecting the right content. Today, many point solutions exist for protecting content, but these are not integrated in to a holistic and comprehensive Secure Content Management strategy. This white paper discusses the challenges and presents considerations to assist you in developing a comprehensive Secure Content Management strategy that will satisfy your company's information security goals and requirements.

## Introduction

The explosive growth in digital content produced by nearly every small to large company requires a true information-centric approach to securing the sensitive information that resides in this content. This white paper defines Secure Content Management and its four key domains, and explains how the security and content management products offered by EMC® Documentum® and RSA, The Security Division of EMC, can be used together to provide the right level of protection, to the right content, at the right time.

## Audience

This white paper is intended for EMC Documentum customers and partners who are interested in securing enterprise content management environments and who would like to learn more about the security capabilities available today in EMC Documentum and RSA product families.

## Defining Secure Content Management

Securing enterprise content management requires a true information-centric security approach that can be structured around four key domains:

### ***Securing access to the content management system***

First, access to the content management system itself needs to be locked down. This requires:

---

Ensuring Security and Compliance of Your EMC Documentum Enterprise Content Management System: A Collaborative Effort of EMC Documentum and RSA

- 
- Integrating the content management system with the enterprise identity management infrastructure (directory, single sign-on)
  - Enforcing strong authentication for user accessing the content management system
  - Preventing password misuse by ensuring single sign-on between the content management system and other web applications such as Microsoft SharePoint.
  - Well-defined access control lists (ACLs) to ensure users, groups, and roles have the right levels of access to the content they need.

### ***Securing content at rest***

Securing content at rest starts with discovering sensitive content within the content management system as well as in file servers and other content servers across the organization. Once sensitive content is identified, it can be protected using the core capabilities of the content management system by setting controls that are aligned with the organization policy:

- Setting ACLs on sensitive content
- Encrypting the content repository and “shredding” content after deletion
- Signing content that requires electronic proof of authenticity and integrity

### ***Securing content in motion and in use***

Securing content at rest is necessary but insufficient. By definition, information is intended to be distributed and used to deliver its business value. A comprehensive approach to securing content management needs to consider securing content in motion and securing content in use.

Multiple controls and technology are available for securing content distributed to and from a content management server:

- Enabling SSL for all communications with the content management system
- Encrypting the most sensitive content and ensuring that only authorized users that receive it can decrypt it for use
- Detecting sensitive content exchanged over the network in violation of established policies. For instance, e-mail to unauthorized recipients

### ***Monitoring and reporting on compliance of the environment***

The last component of a secure enterprise content management deployment is the monitoring of all the controls in place to ensure compliance with the objectives of the security policy. This is achieved by:

- Logging activities in the content management system and on use of the sensitive content across the environment
- Monitoring the network and operating systems in the environment for security events
- Ensuring the systems within the deployment are in-line with the gold configurations specified by corporate policy and detecting and correcting systems that are out of compliance.

The core security capabilities of the Documentum Content Management System augmented with RSA’s security technology provide the essential security capabilities required to secure enterprise content management. The next sections provide more details on the capabilities available today in EMC Documentum and RSA product families.

---

## Securing CMS access

### ***Integrating with the enterprise identity management infrastructure***

The Documentum Content Server (DCS) is designed to integrate seamlessly within the corporate IT infrastructure including enterprise identity management systems. DCS supports connections with multiple directory servers and is integrated with common enterprise directory products including Microsoft Active Directory, Sun ONE Directory Server, and Oracle Internet Directory.

### ***Enforcing strong authentication and enabling single sign-on***

Documentum Content Server implements an open authentication framework that allows the server to support a broad range of strong authentication methods and to easily integrate with customers' enterprise identity management infrastructures. For example, Documentum Content Server leverages its authentication framework to participate in web-based single sign-on (SSO) environments. Out-of-the-box, DCS is configurable to support Web SSO using the RSA Access Manager. DCS is able to authenticate users leveraging multiple strong authentication methods supported by RSA Access Manager including RSA SecurID authentication, digital certificates, and custom authentication methods. In addition, integration with RSA Access Manager allows SSO to be extended to other content management and enterprise resources, such as Microsoft SharePoint servers, which also support the RSA Access Manager (see the RSA guide *RSA SecurBook for Microsoft SharePoint* for more details).

## Securing content at rest

### ***Discovering sensitive content***

The RSA Data Loss Prevention (DLP) Suite is a comprehensive data loss prevention solution that enables customers to discover and protect sensitive data across the enterprise -- within the data center, on the network, and at endpoints. RSA DLP gives customers a better understanding of the location of high-impact business information within their enterprise and enables them to protect against the proliferation of that data.

The RSA DLP Datacenter product, which locates sensitive throughout the data center on file systems, databases, e-mails systems, and large SAN/NAS environments, can be used to scan and identify sensitive information managed by a Documentum Content Server and stored in its associated repositories. Once data is discovered and classified, customers can remediate policy violations and mitigate risks by applying specific data controls to sensitive information such as financial data, intellectual property, or regulated data such as credit card numbers located in the content management environment. Applied data controls can include enablement of the security capabilities available within the Documentum product family (Trusted Content Services and/or Information Rights Management), which are discussed in more detail within this paper.

### ***Controlling access to sensitive content***

Authorization, or entitlement, determines what content can be seen by whom. Documentum assigns authorization at the object level through ACLs, which are automatically applied to objects as they are created. By applying authorization at the object level, every content object, version, and rendition, and every container for content assets from folders to repositories is governed by an ACL throughout its lifecycle.

Robust authorization capabilities ensure that only the appropriate users have access when that access is required. It also reduces the complexity of the enterprise content management system and simplifies navigation by removing from view content that users are not authorized to access.

---

Individual users, groups, or roles can be assigned to ACLs in Documentum. These ACLs determine the level of access a person has to the content. For instance, a person with READ access will only be able to view a document, but not modify it within the repository, whereas a person with VERSION access could edit the document in the repository and replace it with a new version.

Documentum Trusted Content Services (TCS), which are enabled via a license key, provide additional security capabilities that complement the core security features within the Documentum Content Server. TCS provides additional functionality for access in the form of dynamic and logic-based ACLs. These ACLs can be used to define and implement fine-grained access to sensitive content.

For a complete description of all Documentum ACLs and their use, please refer to the white paper *EMC Documentum Security, A Comprehensive Overview*.

## ***Ensuring confidentiality of sensitive data***

### **Repository encryption**

While some organizations choose to rely solely on OS-level security to ensure their content files are secure at rest, those in more secure environments may choose to encrypt their file at rest. Documentum Trusted Content Services (TCS) enables encryption of the file stores that host content assets in the Documentum repository. The encrypted file store prevents access to content files on the operating system level. For example, should intruders compromise OS level security, all they see are encrypted files. This type of security protects against security breaches from the inside, for instance, by an administrator with a malicious intent.

All encryption comes with a performance tradeoff. With repository encryption it will take longer to save and retrieve files since they also have to be encrypted and unencrypted. Encryption can be applied selectively by the file store, allowing for encrypted and unencrypted files in the same repository served by the same Documentum Content Server. This makes it possible to locate some documents in an encrypted store, while locating other, less-sensitive documents in an unencrypted store, thereby avoiding the performance penalty. Additionally, Documentum Content Storage Services provides assignment policies to automate the placement of content among different file stores. This enables sensitive content to be placed in secure file stores based on the content metadata.

The measured performance degradation of Documentum Content Server using encryption is less than 12 percent. Any additional performance degradation caused by audit logs, method executions, and other events, is negligible. The security benefits of TCS, however, are substantial:

- Content security even if OS security is compromised
- Protection against “rogue” administrator
- Secure storage of back-up media offsite
- Secure backup media disposal

Encryption occurs within Documentum Content Server “below the API,” which means that content is exposed through the Documentum API in its unencrypted form. All applications access encrypted content without decoding—as if no encryption were applied. Indexing and full-text search are not impacted by repository encryption. TCS uses the 3DES-CBC encryption algorithm with a 192-bit key length. RSA BSAFE<sup>®</sup> provides the cryptographic functions and algorithm implementations within Documentum TCS.

Repository encryption also applies to any backup conducted at the file system level. As a result, backup media can be safely stored without danger of a security breach. EMC NetWorker<sup>®</sup> for Documentum, for example, enables the backup and restoration of encrypted files from backup tapes containing encrypted data.

Repository encryption is applied only to content files, not to its metadata (content properties). Nevertheless, since metadata is kept in a standard relational database, any RDBMS security measure provided by the database vendor can be employed. Oracle, for example, provides a complete database encryption solution.

---

For those in a distributed environment using Documentum Branch Office Caching Services (BOCS), encryption is also available for the remote content caches. This feature is a standard part of BOCS and does not require the installation of TCS.

## **Digital shredding**

Many end users do not realize that when they delete their files, the files are not really gone. Deleted files may be recovered by analyzing the magnetic traces data leaves behind on its storage medium. This is how the FBI retrieves data from subpoenaed computers.

Digital shredding permanently destroys file data when the OS delete/unlink command is issued. It automatically writes over the location of data multiple times to ensure that it cannot be recovered, even by analyzing residual magnetism. The default setting for digital shredding in Documentum Trusted Content Services (TCS) is three overwrites; however, this is configurable and may be set higher if deemed necessary. This feature of TCS supports records management and retention policy applications that define when in its lifecycle content should be disposed. Digital shredding is considered a mandatory step for most records management applications.

## ***Ensuring authenticity and integrity of sensitive data***

Documentum Trusted Content Services (TCS) enables any content asset or event such as a business process task to be signed electronically. Electronic signatures are securely linked to the content object and stored in the Documentum repository as part of the audit trail. Any subsequent modification of the content invalidates the stored signature.

Signatures are date and time-stamped and include the name of the person signing along with a justification for that signature. Each signature also contains a hash-checksum that verifies the authenticity of the signed content. Only valid signatories are permitted to enter an electronic signature.

Documentum also provides the foundation for legally admissible digital signatures. A digital signature is the data element that allows the recipient of a message or transaction to verify its content and signatory. It's the electronic equivalent of having a paper document signed and notarized and it is not a digitized image of a handwritten signature. Users are validated using strong authentication (that is, PKI cryptography) instead of just a username/password pair. Digital signatures are portable and can be verified outside of the signer's organization or location.

## **Security content in motion and in use**

### ***Enabling secure communication***

In addition to encrypting content at rest, all data traffic can be encrypted between Documentum Content Server and Documentum clients, including Documentum CenterStage<sup>®</sup>, using the secure sockets layer (SSL) standard. Additionally, data traffic between Documentum Content Server and directory servers can be encrypted using SSL. Documentum Content Server can be set up to use secured or unsecured ports and clients can be mandated to connect through secured ports only. Documentum SSL employs the ADH (Anonymous Diffie-Hellman) algorithm with a 1024-bit key size for key exchange. Data encryption uses the AES algorithm with 256-bit keys.

Encrypted communication prevents eavesdropping and ensures data privacy. It increases the flexibility of network architecture and allows directory servers or Documentum Content Server to be placed inside or outside a DMZ.

### ***Enforcing sensitive content policies***

Products from the RSA DLP Suite can be used by customers to locate, monitor, and control data in motion and in use within the content management environment. The RSA DLP Network product provides

---

monitoring and enforcement capabilities for transmission of sensitive data on a network including e-mail and web traffic. For example, DLP technology can be integrated with an enterprise's e-mail gateway to monitor the exchange of sensitive data via e-mail and to enforce protection (for example, encryption) or quarantine policies on e-mails containing sensitive data.

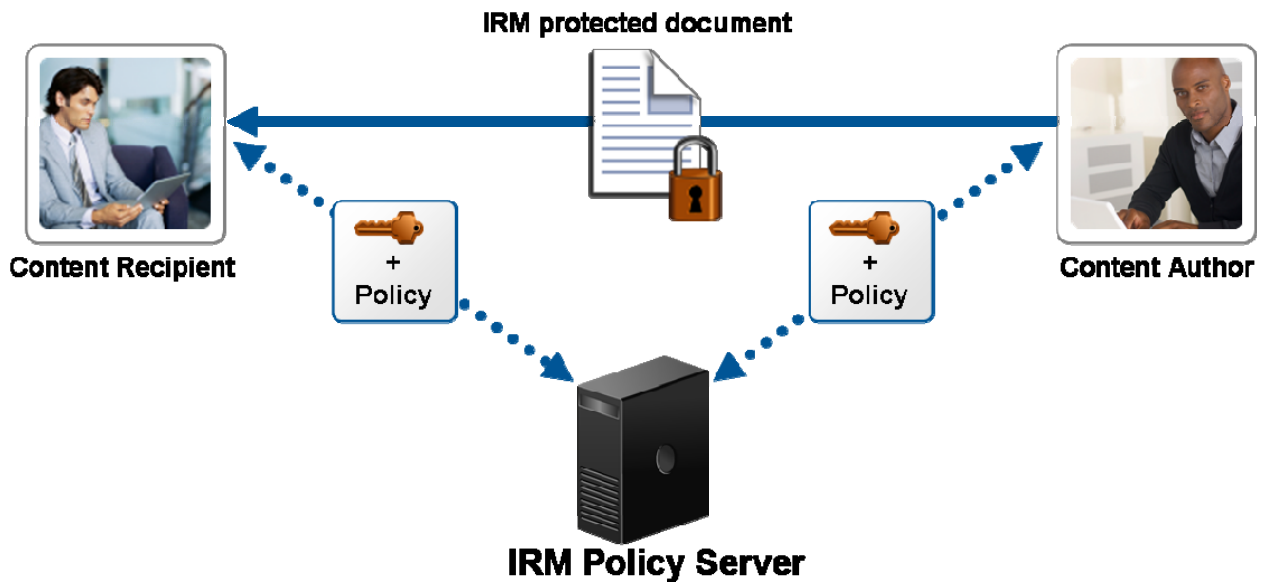
The RSA DLP Endpoint product provides discovery, monitoring, and enforcement capabilities for sensitive data residing on endpoints such as desktops and laptops.

## ***Information Rights Management***

While the Documentum repository does an excellent job of protecting content at rest and ensuring only those authorized can access content, content must often travel outside the repository and beyond the firewall in the normal course of business. The enterprise must reach beyond itself to enable vendors, distributors, contractors, regulators, and other partners to share content across the entire value chain. When dealing with sensitive information, organizations need to control how content is viewed, printed, copied, and edit, and by whom, regardless of where it resides. Ensuring security of content once it leaves the repository is a challenge. Even the most confidential of documents may be compromised once they are opened on a desktop, e-mailed, or printed. That's where EMC Documentum Information Rights Management (IRM) comes in. IRM can persistently protect content when it leaves the repository, even if it travels beyond the firewall.

The Documentum IRM Server, and the various client components, form a system with the primary objective of providing strong, persistent security and control over information in electronic form outside the repository. IRM advances the security of your Documentum repository in two fundamental ways. First, it extends the access controls governing content inside the repository to follow the content wherever it goes outside the repository. Second, it extends the types of permissions that can be enforced on content to include control over the ability to print, copy, modify local copies, apply visual watermarks and prevent screen capture. It accomplishes this through the use of strong cryptography, controlling access to decryption keys and locking down viewing applications to prevent undesired use of the information once decrypted.

The Documentum IRM suite of products consist primarily of a Policy Server and client plug-ins that enable certain business applications, such as Microsoft Office and Adobe Acrobat, to view and protect documents with IRM. A document or e-mail that is deemed confidential can be protected with IRM either programmatically by an application or by the content author. The content is encrypted and assigned policies governing its usage. Both the encryption key and the policy are stored on the IRM Policy Server. Once a document or e-mail is encrypted with IRM, it can safely reside anywhere. To open protected content, the recipient uses the same business application they normally would, like Microsoft Excel or Outlook for instance. The corresponding client plug-in will contact the IRM Policy Server and once the user authenticates, will be granted the appropriate privileges to the content.



**Figure 1. Accessing IRM-protected content**

Since the IRM policy and the encryption key are stored on the policy server, separate from the content, it doesn't matter where the content goes. Even if copies are stored on fixed media, such as CD or DVD, the content author still has the ability to change or revoke the privileges on the policy server at any time. Expiring the policy will render all copies useless. The author is always in control of the content.

RSA BSAFE® provides the strong, FIPS compliant 256-bit AES cryptography utilized within IRM. This implementation is secure enough to protect sensitive federal government documents. It is also commonly used to protect intellectual property during collaboration and legal communications and to protect personal and customer data in compliance with regulations such as HIPAA and Mass 201 CMR 17.

For a complete description of Documentum IRM capabilities please refer to the white paper *EMC Documentum Information Rights Management, Overview of Technical Architecture*.

## Compliance

### ***Logging in the content management system***

Logging is an important part of any secure system where auditing and compliance are a concern. Logging in Documentum Enterprise Content Management can be tuned to provide the right amount of logging need for any given deployment. Organizations can determine which events they want to log and how much data to capture. Applications built on top of Documentum can likewise add any custom events to the Documentum logs.

Logs, or audit trails, can be exposed through any number of reporting mechanisms and notifications can provide the desired amount of system monitoring. Notifications can, for instance, send an e-mail to a system administrator if someone tries to print an IRM-protected document more than twice. An audit report may contain information such as the user ID, date/time, and IP address of everyone who has retrieved a classified document from the repository.

The audit trails themselves are kept secure and give organizations the ability to prove an audit trail has not been altered.

---

## **Monitoring and reporting security events**

RSA enVision is a security and information event management (SIEM) and log management solution. RSA enVision provides collection, alerting, and analysis of log data that enables organizations to simply comply and quickly respond to high-risk security events.

RSA enVision can be used to monitor the log and audit information produced by the host platforms and the devices that comprise the content management environment. In addition, RSA enVision can collect logs from the RSA DLP Datacenter product that can be used to scan the Documentum Content Server. This allows enVision to correlate information about the creation and location of sensitive data with information about server vulnerability and anomalous user activity. This correlation provides security administrators with a more holistic view of the content management environment and enables them to better prioritize security vulnerability response and remediation.

## **Conclusion**

Securing an information infrastructure is all about ensuring that the right people have access to the right information over a trusted infrastructure and using a process that is easy and efficient to manage. EMC Documentum Enterprise Content Management together with RSA technology provides the core capabilities to:

- Discover sensitive content,
- Define and enforce policies governing the protection of sensitive content at rest, in motion and in use, and
- Monitor the environment to ensure compliance to policies.

## **References**

For further information on the RSA products discussed in this paper please see <http://www.rsa.com/node.aspx?id=1155>.

For further information on the Documentum Content Management products discussed in this paper please see <http://www.emc.com/products/category/content-management.htm>.

*EMC Documentum Security, A Comprehensive Overview* — [http://info.emc.com/mk/get/SDL?reg\\_src=WEB&P.ctp\\_program\\_execution.Source\\_ID=6139](http://info.emc.com/mk/get/SDL?reg_src=WEB&P.ctp_program_execution.Source_ID=6139)

*EMC Documentum Information Rights Management, Overview of Technical Architecture* — [http://info.emc.com/mk/get/SDL?reg\\_src=web&P.ctp\\_program\\_execution.Source\\_ID=15308](http://info.emc.com/mk/get/SDL?reg_src=web&P.ctp_program_execution.Source_ID=15308)

RSA Solution for Microsoft SharePoint — <http://www.rsa.com/node.aspx?id=1333>